Article

The International Journal of Robotics Research
2018, Vol. 37(2–3) 249–265

DOI: 10.1177/0278364918755536
journals.sagepub.com/home/ijr

A convex polynomial model for planar sliding mechanics: theory, application, and experimental validation

Jiaji Zhou, Matthew T Mason, Robert Paolini and Drew Bagnell

Abstract

We propose a polynomial model for planar sliding mechanics. For the force–motion mapping, we treat the set of generalized friction loads as the 1-sublevel set of a polynomial whose gradient directions correspond to generalized velocities. The polynomial is confined to be convex even-degree homogeneous in order to obey the maximum work inequality, symmetry, shape invariance in scale, and fast invertibility. We present a simple and statistically efficient model identification procedure using a sum-of-squares convex relaxation. We then derive the kinematic contact model that resolves the contact modes and instantaneous object motion given a position controlled manipulator action. The inherently stochastic object-to-surface friction distributions are modeled by sampling polynomial parameters from distributions that preserve sum-of-squares convexity. Thanks to the model smoothness, the mechanics of patch contact is captured while being computationally efficient without mode selection at support points. Simulation and robotic experiments on pushing and grasping validate the accuracy and efficiency of our approach.

Keywords

Mechanics of manipulation, frictional contact modeling, sum-of-squares method

1. Introduction

Effective robotic manipulation requires an understanding of the underlying physical processes. Mason (1986a) explored using pushing as a sensorless mechanical funnel to reduce uncertainty. Whitney (1983) analyzed the mechanics of wedging and jamming during peg-in-hole insertion and designed the Remote Center Compliance device that significantly increases the success of the operation under motion uncertainty. With a well-defined generalized damper model, Lozano-Perez et al. (1984) and Erdmann (1986) developed strategies to chain a sequence of operations, each with a certain funnel, to guarantee operation success despite uncertainty. These successes stem from robustness analysis using simple physics models.

Planning and control without explicit reasoning about uncertainty and the task mechanics can lead to undesirable results. For example, grasp planning (Ferrari and Canny, 1992; Miller et al., 2003) is often prone to failure: the object moves while the fingers close and ends up in a final relative pose that differs from planned. Consider the process of closing a parallel jaw gripper shown in Figure 1, the object will slide when the first finger engages contact and pushes the object before the other one touches the object. If the object does not end up slipping out, it can be jammed at an undesirable position or grasped at an unexpected position. A high-fidelity and easily identifiable model with minimum adjustable parameters capturing all these possible outcomes would enable the synthesis of a robust manipulation strategy.

We develop a data-driven but physics-based method for modeling planar friction. Manipulations employing friction are ubiquitous in tasks including positioning and orienting objects by pushing (Akella and Mason, 1998; Dogar and Srinivasa, 2010; Lynch and Mason, 1996; Mason, 1986a), controlled slip with dexterous hands (Cole et al., 1992) and assembly of tight-fitting parts (Whitney, 1983). For planar manipulation with finite object motion, indeterminacy of the pressure distribution between the object and support surface leads to uncertainty in the resultant velocity. Despite such inherent difficulty, algorithms and analyses have been developed with provable guarantees. Mason (1986a) derived the voting theorem to determine the sense of rotation of an object pushed by a point contact. Lynch and Mason (1996) developed a stable pushing strategy when objects remain fixed to the end effector with two or more contact points. However, minimal assumptions on friction

Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, USA

Corresponding author:
Jiaji Zhou, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA.
Email: jiajiz@cs.cmu.edu

conditions inherently lead to conservative strategies. By explicitly modeling and identifying the friction space, we can improve strategies for planning and control.

The paper combines and extends our previous work (Zhou et al., 2017, 2016). The first contribution is the development of a precise and statistically efficient force–motion model with a computationally efficient identification procedure. We propose a framework representing planar sliding force–motion models using homogeneous even-degree *sum-of-squares* convex (sos-convex) polynomials, which can be identified by solving a semi-definite programming. The set of applied friction wrenches is the 1-sublevel set of a convex polynomial whose gradient directions correspond to incurred sliding body twist. The second contribution is a quasi-static kinematic contact solution for manipulation problems with finite planar sliding motion. The algorithm maps a commanded rigid position-controlled end-effector motion to the instantaneous resultant object motion, with detection of the equilibrium state (jamming or grasping). The applied wrench is solved as an intermediate output. We show that single contact with a convex quadratic force–motion model has a unique analytic linear solution that extends the approach in Lynch et al. (1992). The case for a high-order convex polynomial force–motion model is reduced to solving a sequence of such subproblems. For multiple contacts (e.g. pushing with multiple points or grasping) we need to add linear complementarity constraints (Stewart and Trinkle, 1996) at the pusher points, and the entire problem is a standard linear complementarity problem (LCP).

The inherent stochasticity in frictional sliding is modeled by sampling the physics parameters from proper distributions. We validate the model by comparing the simulation with large-scale experimental data on robotic pushing and grasping tasks. The model serves as a good basis for both open-loop planning and feedback control.[1]

We assume quasi-static rigid-body planar mechanics (Mason, 1986b) where inertia forces and out-of-plane moments are negligible. Figure 2 illustrates the outline of this article. The rest is organized as follows.

- Section 2 describes the previous work.
- Section 3.1 reviews the background of force–motion model for sliding.
- Section 3.2 develops the convex polynomial representation and the identification algorithm.
- Section 3.3 demonstrates model identification results based on simulation and experimental data.
- Section 3.4 demonstrates stable push action generation based on the invertible property of the model.
- Section 4 develops the kinematic contact model for single and multiple frictional contacts.
- Section 4.3 develops the sampling strategies of physically consistent model parameters that captures the inherent frictional stochasticity.
- Sections 4.4–4.6 demonstrates simulation and experimental results for pushing and grasping applications.

2. Related work

The mechanics of pushing and grasping involving finite object motion with frictional support was first studied in Mason (1986a). A notable result is the voting theorem which dictates the sense of rotation given a push action and the center of pressure regardless of the uncertain pressure distribution. Brost (1988) used this result to construct the operational space for planning squeezing and push-grasping actions under uncertainty. However, many unrealistic assumptions are made in order to reduce the state space and create finite discrete transitions, including infinitely long fingers approaching the object from infinitely far away. In addition, how far to push the object in the push-grasp action was not addressed. Peshkin and Sanderson (1988a) provided an analysis on the slowest speed of rotation given a single point push. Peshkin and Sanderson (1988b) used this result to design fences for parts feeding. Lynch and Mason (1996) derived conditions for stable edge pushing such that the object will remain attached to the pusher without slipping or breaking contact. All these results do not assume knowledge of the pressure distribution except the location of the center of pressure. They can be classified as worse-case guarantees without looking into the details of sliding motion. Despite being agnostic to pressure distribution, these methods tend to be overly conservative, have impractical assumptions, or both.

Friction parameter estimation has been proposed to improve planning and control. Yoshikawa and Kurisu (1991) solved an unconstrained least-squares problem to estimate the center of friction and the pressure distribution over discrete grids on the contact surface. With a similar set up, Lynch (1993) proposed a constrained linear programming procedure to avoid negative pressure assignment. However, methods based on discretization of the support surface introduce two sources of error in both localization of support points and pressure assignment among those points. We do not need to estimate the exact location of support points. Coarse discretization loses accuracy while fine discretization suffers from the curse of dimensionality. In addition, the instantaneous center of rotation (COR) of the object can coincide with one of the support points, rendering the kinematic solution computationally hard due to combinatorial sliding/sticking mode assignment for each support point.

Goyal et al. (1991) noted that all the possible static and sliding frictional wrenches, regardless of the pressure distribution, form a convex set whose boundary is called a limit surface (LS). Analytic construction of a LS from the Minkowski sum of frictional limit curves at individual support points, however, is intractable. Howe and Cutkosky (1996) presented an ellipsoid approximation of the LS assuming known pressure distribution. The ellipsoid is constructed by computing or measuring the major axis lengths (maximum force during pure translation and maximum torque during pure rotation). Facets can be added by intersecting the ellipsoid with planes determined by each

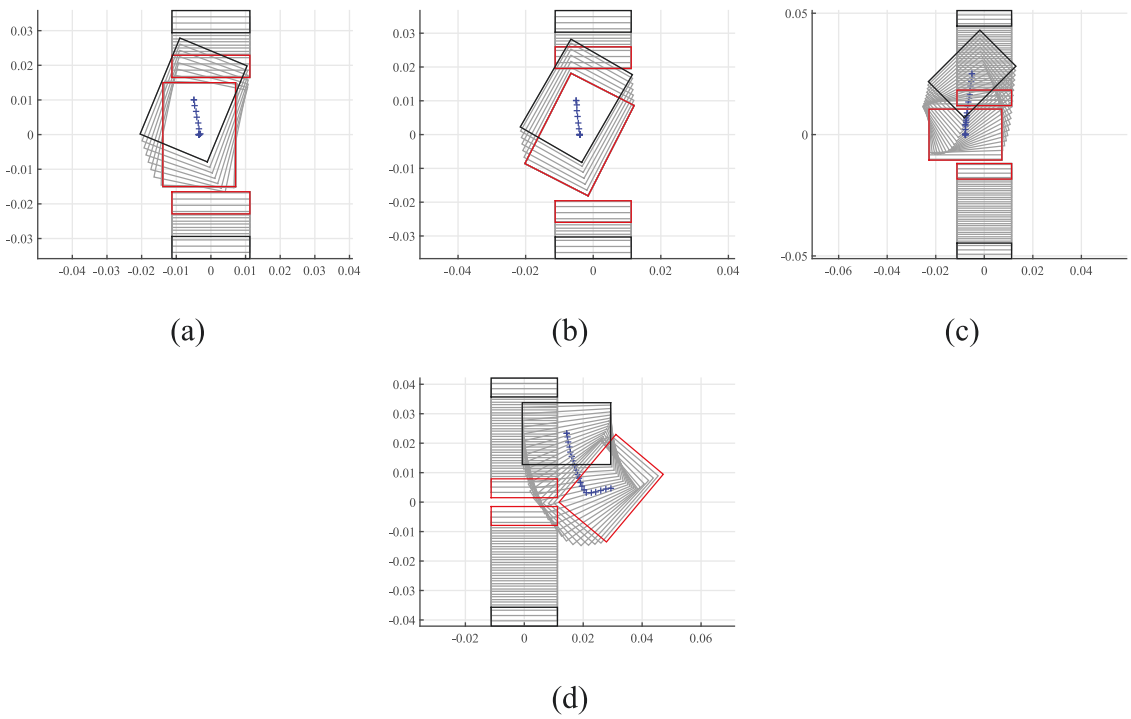

Fig. 1. Simulation results using the proposed contact model illustrating the process of a parallel jaw gripper squeezing along the y axis when the object is placed at different initial poses. The initial, final, and intermediate gripper configurations and object poses are shown in black, red, and gray, respectively. Blue plus signs trace out the center of mass (COM) trajectory of the object. (a) Grasped with offset. (b) Jamming. (c) Grasped with offset. (d) Slipped to free space.

support point. The pressure distribution (except for computable three-point support with known center of pressure), nevertheless, is non-trivial to measure. We also show that the ellipsoid approximation, as the convex quadratic special case of our convex polynomial representation, is less accurate due to the lack of expressiveness, particularly when the support regions are scattered.

Lynch et al. (1992) derived the kinematics of single-point pushing with a centered and axis-aligned ellipsoid approximation. Based on the kinematic model, Hogan and Rodriguez (2016) recently proposed using hybrid model predictive control to generate a sequence of open-loop pushing commands. Yu et al. (2016) recently presented a large-scale empirical effort to verify the generalized Coulomb friction law (Moreau, 1988) and demonstrated inherent stochasticity in planar sliding motions.

Recent data-driven attempts (Kopicki et al., 2011; Omrcen et al., 2009) collected visual data from random push trials and applied "off-the-shelf" machine learning algorithms to build forward motion models. Our model identification also embraces a data-driven strategy but bears in mind that physics principles should guide the design of the learning algorithm (as constraints and/or priors), hence increasing data efficiency and generalization performance. For example, our training procedure only requires a few data points (less than 10) and a change of uniform surface material or a scaling in object mass does not require retraining. In addition, different applications based on the model, as shown later, are physically consistent and computationally efficient.

3. Force–motion model

We first introduce the following notation.

- O: the object center of mass (COM) used as the origin of the body frame.[2] We assume vector quantities are with respect to the body frame unless specifically noted.
- R: the region between the object and the supporting surface.
- $\mathbf{f_r}$: the surface friction force applied by the object at a point $\mathbf{r}$ in R.
- $\mathbf{V} = [V_x; V_y; \omega]$: the body twist (generalized velocity).
- $\mathbf{F} = [F_x; F_y; \tau]$: the generalized friction load. Here $\mathbf{F}$ equals the applied body wrench by the manipulator when the object is in quasi-static balance.
- $\mathbf{p}_i$: each contact point between the manipulator end effector and object in the body frame.
- $\mathbf{v}_{p_i}$: applied velocities by the manipulator end effector at each contact point in the body frame.
- $\mathbf{n}_{p_i}$: the inward normal at contact point $\mathbf{p}_i$ on the object.
- μ_c: coefficient of friction between the object and the manipulator end effector.

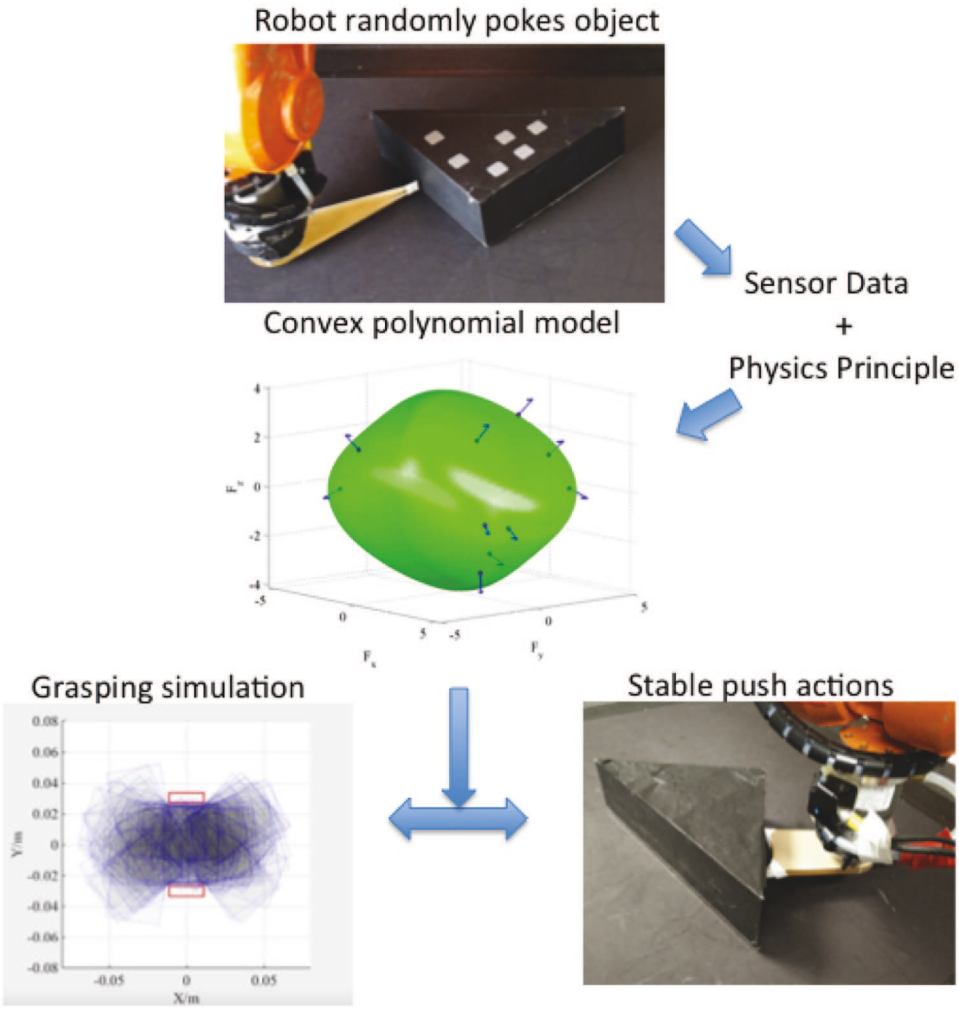

Fig. 2. The robot randomly pokes the object of known shape with a point finger to collect force–motion data. We then optimize a convex polynomial friction representation with physics-based constraints. Stable pushing and grasping simulations under pose uncertainty are two example applications of the model.

3.1. Background on planar friction

The classical Coulomb friction law states that for a point contact with instantaneous planar velocity $\mathbf{v} = [v_x, v_y]^{\mathrm{T}}$, the incurred friction force $\mathbf{f} = [f_x, f_y]^{\mathrm{T}}$ the point applies on the surface is parallel to $\mathbf{v}$, i.e. $\mathbf{f}/|\mathbf{f}| = \mathbf{v}/|\mathbf{v}|$. We refer the readers to Mason (1986a) for details of the friction analysis for planar sliding under the isotropic Coulomb friction law. In this paper, we build our analysis on a generalized friction law formulated first in Moreau (1988), in which $\mathbf{v}$ and $\mathbf{f}$ may not be parallel, but only need to obey the maximum work inequality:

$$(\mathbf{f} - \mathbf{f}') \cdot \mathbf{v} \geq 0 \tag{1}$$

where $\mathbf{f}'$ is an arbitrary element from the set of all possible static and sliding friction forces.

We can compute the generalized friction load $\mathbf{F}$ by integration over R:

$$F_x = \int_R f_{rx}\, dr, F_y = \int_R f_{ry}\, dr, \tau = \int_R (r_x f_{ry} - r_y f_{rx})\, dr \tag{2}$$

The maximum work inequality in Equation (1) can be extended to the generalized friction load $\mathbf{F}$ and twist $\mathbf{V}$:

$$\begin{aligned} \mathbf{F} \cdot \mathbf{V} &= \int_R f_{rx}(V_x - \omega r_y)\, dr + \int_R f_{ry}(V_y + \omega r_x)\, dr \\ &= \int_R f_{rx} v_{rx} + f_{ry} v_{ry}\, dr = \int_R \mathbf{f_r} \cdot \mathbf{v_r}\, dr \geq \mathbf{F}' \cdot \mathbf{V} \end{aligned} \tag{3}$$

among any other possible generalized friction load $\mathbf{F}'$. Owing to the converse supporting hyperplane theorem (Boyd and Vandenberghe, 2004), the set of all generalized friction loads form a convex set $\mathcal{F}$. An important work that inspires us is Goyal et al. (1991) who found that all possible generalized friction loads during sliding form a LS constructed from the Minkowski sum of limit curves at individual support points. Points inside the surface correspond to static friction loads. Points on the surface correspond to friction loads with normals parallel to sliding velocity directions, forming a mapping between the generalized friction load and sliding velocity. An ideal LS is always convex due

to the maximum work inequality but may not be strictly convex when a single point supports finite pressure. As shown in Figure 4b, facets can occur since the object can rotate about one of the three support points whose velocity is zero with indeterminate underlying friction.

Erdmann (1994) proposed a configuration space embedding of friction. In this work, the third component of $\mathbf{F}$ is $F_z = \tau/\rho$ and the third component of $\mathbf{V}$ is $V_z = \omega\rho$, where ρ is the radius of gyration. In doing so, all three components in $\mathbf{F}$ and $\mathbf{V}$ have the same unit. Observe that such a normalized representation also obeys maximum work inequality with ρ being any characteristic length. In our experiments, we have found that the normalized representation yields better numerical condition and different values of ρ including radius of gyration, average edge length, and minimum enclosing circle radius lead to similar performance.

3.2. Representation and identification

In this section, we propose the sublevel set representation of friction with desired properties and show that convex even-degree homogeneous polynomials are valid solutions. Then we formulate an efficient convex optimization procedure to identify such polynomials.

3.2.1. Polynomial sublevel set representation. Let $H(\mathbf{F})$ be a differentiable convex function that models the generalized friction load and velocity as follows.

- The 1-sublevel set $L_1^-(H) = \{\mathbf{F} : H(\mathbf{F}) \leq 1\}$ corresponds to the convex set $\mathcal{F}$ of all generalized friction loads.
- The 1-level set $L_1(H) = \{\mathbf{F} : H(\mathbf{F}) = 1\}$ corresponds to generalized friction loads (during slip) on the boundary surface of $\mathcal{F}$.
- The surface normals given by gradients $\{\nabla H(\mathbf{F}) : \mathbf{F} \in L_1(H)\}$ represent instantaneous generalized velocity directions during slip, i.e. $V = s\nabla H(\mathbf{F})$ where $s > 0$.

Theorem 1. *The set of friction loads represented by the* 1*-sublevel set of a differentiable convex function follows the maximum work inequality.*

Proof. When the object remains static, $\mathbf{F}$ belongs to the interior of $L_1^-(H)$, and $\mathbf{V}$ equals zero, the inequality holds as equality. When the object slips, $\mathbf{F} \in L_1(H)$, and $\mathbf{V}$ is non-zero, we have for any other generalized friction load $\mathbf{F}' \in L_1^-(H)$:

$$\mathbf{V}\cdot(\mathbf{F}' - \mathbf{F}) = s(\nabla H(\mathbf{F})\cdot(\mathbf{F}' - \mathbf{F})) \leq s(H(\mathbf{F}') - H(\mathbf{F})) \leq 0$$

where the first inequality is due to the convexity of $H(\mathbf{F})$. ■

In addition to enforcing convexity (discussed in 3.2.2), we choose $H(\mathbf{F})$ to obey the following properties.

1. Symmetry: $H(\mathbf{F}) = H(-\mathbf{F})$ and $\nabla H(\mathbf{F}) = -\nabla H(-\mathbf{F})$.
2. Scale invariance: $\nabla H(s\mathbf{F}) = g(s)\nabla H(\mathbf{F})$, where $g(s)$ is a positive scalar function over scalar s.
3. Efficient invertibility: efficient numerical procedures exist for finding a $\mathbf{F} \in L_1(H)$ such that $\nabla H(\mathbf{F})/\|\nabla H(\mathbf{F})\| = \mathbf{V}$ for a given query unit velocity $\mathbf{V}$. We denote such an operation as $\mathbf{F} = H_{inv}(\mathbf{V})$.

Symmetry is based on the assumption that negating the velocity direction would only result in a sign change in the friction load. Scale invariance is desirable for two reasons: (1) scaling in mass and surface coefficient of friction could only result in a change of scale but not other geometrical properties of the level-set representation; and (2) predicting the directions of generalized velocities (by computing gradients and normalizing to a unit vector) only depends on the direction of generalized force. Such a property is useful in the context of pushing with robot fingers where applied loads are represented by friction cones. The inverse problem of finding the friction load for a given velocity naturally appears in seeking quasi-static balance for stable pushing or computing deceleration during free sliding. In general, an efficient numerical solution to the inverse problem, which our representation enables, is key to planning and simulation. One solution family for $H(\mathbf{F})$ that obeys these properties is the set of *strictly convex even-degree homogeneous polynomials*.

Theorem 2. *A strictly convex even degree-d homogeneous polynomial* $H(\mathbf{F}; a) = \sum_{i=1}^{m} a_i (F_x)^{i_1} (F_y)^{i_2} (F_z)^{d-i_1-i_2}$ *with m (bounded by* $\binom{d+2}{2}$*) monomial terms parametrized by a satisfies the properties of symmetry, scale invariance, and efficient invertibility.*

Proof. Proving symmetry and scale invariance are trivial due to the homogeneous and even-degree form of $H(\mathbf{F})$. Here, we sketch the proof that efficient invertibility can be achieved by first solving a simple non-linear least-squares problem followed by a rescaling.

Construct an objective function $G(\mathbf{F}) = \frac{1}{2}\|\nabla H(\mathbf{F}) - \mathbf{V}\|^2$ whose gradient $\frac{\partial G}{\partial \mathbf{F}} = \nabla^2 H(\mathbf{F})(\nabla H(\mathbf{F}) - \mathbf{V})$. Note that its stationary point F^*, which iterative methods such as the Gauss–Newton or trust-region algorithms will converge to, satisfies $\nabla H(\mathbf{F}^*) - \mathbf{V} = 0$. Hence, F^* is globally optimal with value zero. Let $\Delta\mathbf{F}_t = \nabla^2 H(\mathbf{F}_t)^{-1}(\mathbf{V}_t - \mathbf{V})$,[3] then the update rule for the Gauss–Newton algorithm is $\mathbf{F}_{t+1} = \mathbf{F}_t - \Delta\mathbf{F}_t$. Although the final iteration point $\mathbf{F}_T$ may not lie on the 1-level set of $H(\mathbf{F})$, we can scale $\mathbf{F}_T$ by $\hat{\mathbf{F}}_T = H(\mathbf{F}_T)^{-1/d}\mathbf{F}_T$ such that $H(\hat{\mathbf{F}}_T) = 1$ and $\nabla H(\hat{\mathbf{F}}_T)/\|\nabla H(\hat{\mathbf{F}}_T)\| = V$ due to the homogeneous form of $H(\mathbf{F})$. Therefore, $H_{inv}(\mathbf{V}) = \hat{\mathbf{F}}_T$. ■

3.2.2. Sum-of-squares convex relaxation. Enforcing strong convexity for a degree-two homogeneous polynomial $H(\mathbf{F}; A) = \mathbf{F}^{\mathrm{T}} A\mathbf{F}$ has a straightforward set up as solving a semi-definite programming problem with constraint of $A \succeq \epsilon I$. Meanwhile, for a polynomial of degree greater than three whose Hessian matrix $\nabla^2 H(\mathbf{F}; a)$

is a function of both $\mathbf{F}$ and a, certification of positive semi-definiteness is NP-hard. However, recent progress (Magnani et al., 2005; Parrilo, 2000) in sum-of-squares programming has given powerful semi-definite relaxations of global positiveness certification of polynomials. Specifically, let $\mathbf{z}$ be an arbitrary non-zero vector in $\mathbb{R}^3$ and $y(\mathbf{F},\mathbf{z}) = [z_1F_x, z_1F_y, z_1F_z, z_2F_x, z_2F_y, z_2F_z, z_3F_x, z_3F_y, z_3F_z]^{\mathrm{T}}$. If there exists a positive-definite matrix Q such that

$$\mathbf{z}^{\mathrm{T}}\nabla^2 H(\mathbf{F};a)\,\mathbf{z} = y(\mathbf{F},\mathbf{z})^{\mathrm{T}}\, Q y(\mathbf{F},\mathbf{z}) > 0 \tag{4}$$

then $\nabla^2 H(\mathbf{F};a)$ is positive definite for all non-zero $\mathbf{F}$ under parameter a and $H(\mathbf{F};a)$ is called sos-convex. Further, Equation (4) can be written as a set of K sparse linear constraints on Q and a.

$$\begin{aligned} &\mathrm{Tr}(A_k Q) = b_k^{\mathrm{T}} a, \quad k \in \{1 \dots K\} \\ &Q - \epsilon I \succeq \end{aligned} \tag{5}$$

where A_k and b_k are a constant sparse element indicator matrix and vector that only depend on the polynomial degree d. The symbol $\succeq$ implies that the left-hand side is positive semidefinite. The number of constraints K equals 27 for $d = 4$.

3.2.3. Identification. This section sets up an efficient convex optimization for identifying the coefficient a of the polynomial $H(\mathbf{F};a)$ given a set of measured noisy generalized force–motion $\{\mathbf{F}_{i\in\{1\dots N\}}, \mathbf{V}_{i\in\{1\dots N\}}\}$ pairs. In our experiments, we use a homogeneous fourth-order polynomial. The optimization should find the coefficient a such that the measured forces $\mathbf{F}_i$ are close to the 1-level set surface and the corresponding gradients are aligned well (up to scale) with respect to measured velocities $\mathbf{V}_i$. Let $\alpha_i = \|\nabla H(\mathbf{F}_i;a) - (\nabla H(\mathbf{F}_i;a)\cdot V_i)\, V_i\|_2^2$ be the L2-projection residual of $\nabla H(\mathbf{F}_i;a)$ onto the measured unit velocity vector V_i, and let $\beta_i = (H(\mathbf{F}_i;a) - 1)^2$ be a distance measurement of $\mathbf{F}_i$ from the 1-level set of $H(\mathbf{F}_i;a)$. We set up the optimization as follows:

$$\underset{a,Q}{\text{minimize}} \quad \|a\|_2^2 + \sum_{i=1}^{N} (\eta_1\alpha_i + \eta_2\beta_i) \tag{6}$$

$$\text{subject to} \quad \mathrm{Tr}(A_k Q) = b_k^{\mathrm{T}} a, \quad k = 1, \dots, K \tag{7}$$

$$Q - \epsilon I \succeq 0 \tag{8}$$

The first term is for parameter regularization. η_1 and η_2 are trade-off parameters determined by cross-validation. Equations (7) and (8) enforce convexity. Note that the objective is quadratic in a with sparse linear constraints and a semi-definite constraint on Q. We would like to point out that the formulation can be adapted online using projected gradient descent so that the importance of historical data is diminishing as the object moves, enabling the estimation to adapt to changing surface conditions.

3.3. Identification experiments

We conduct simulation and robotic experiments to demonstrate the accuracy and statistical-efficiency of our proposed representation. The model converges to a good solution with few available data, which saves experimental time and design efforts. We compare the following four different force–motion model representations $\mathcal{H}$: (1) degree-four convex homogeneous polynomial (poly4-cvx); (2) degree-four homogeneous polynomial (poly4) without convexity constraints (3) convex quadratic (quad) as a degree-two polynomial, i.e. $H(\mathbf{F}) = \mathbf{F}^{\mathrm{T}} A \mathbf{F}$ with ellipsoid sublevel set; and (4) Gaussian process (GP) with squared exponential kernel as a smooth generic non-parametric data efficient learning model.[4]

Denote by $\mathbf{V}_i$ the ground truth instantaneous generalized velocity direction and $\mathbf{V}_p(\mathbf{F}_i;\mathcal{H})$ as the predicted generalized velocity direction based on $\mathcal{H}$ for the input generalized load $\mathbf{F}_i$, we use the average angle $\delta(\mathcal{H}) = \frac{1}{N}\sum_{i=1}^{N} \arccos(\mathbf{V}_p(F_i;\mathcal{H})\cdot\mathbf{V}_i)$ between $\mathbf{V}_p(F_i;\mathcal{H})$ and $\mathbf{V}_i$ as an evaluation criterion.

3.3.1. Simulation study. Two kinds of pressure distribution are studied.

- "Legged" support: Randomly sampled three support points on a unit circle with randomly assigned pressure.
- "Uniform" support: Uniformly distributed 360 support points on a unit circle and 400 support points within a unit square. Each point has the same support pressure.

For each pressure configuration, we conduct 50 experimental trials. To generate the simulated force–motion data, we assume a Coulomb friction model at each support point with a uniform coefficient of friction. Without loss of generality, sum of pressure over all contact points is normalized to one and the origin is set as the center of pressure. For each trial of "uniform" support, we sampled 150 instantaneous generalized velocities directions $\mathbf{V}_i$ uniformly on the unit sphere and compute the corresponding generalized friction loads $\mathbf{F}_i$. For each trial of "legged" support, 75 $(\mathbf{F}_i, \mathbf{V}_i)$ pairs are uniformly sampled on the facets (same $\mathbf{V}_i$ but different $\mathbf{F}_i$ for each facet) and another 75 pairs are uniformly sampled in the same fashion as "uniform" support. In doing so, the dataset has a diverse coverage. Among the 150 pairs, 50% is used for hold-out testing, 20% is used for cross-validation and four different amounts (7, 15, 22, 45) from the remaining 30% are used as training. In order to evaluate the algorithms' robustness under noise, we additionally corrupt the training and validation set using Gaussian noise of standard deviation $\sigma = 0.1$ to each dimension of both $\mathbf{F}_i$ and $\mathbf{V}_i$ (renormalized to unit vector). From Figure 3 we can draw the following conclusions. (1) Poly4-cvx has the smallest $\delta(\mathcal{H})$ for different amounts of training data and pressure configurations. (2) Both poly4-cvx and convex quadratic show superior performance when data is scarce and noisy, demonstrating convexity is key to

data-efficiency and robustness. Poly4-cvx model additionally shows larger improvement as more data is available due to stronger model expressiveness. (3) Poly4 (without convexity constraint) performs the worst when only few data is available, but gradually improves as more data is available for shaping the surface. For noise-free experiments shown in Figure 3b and d, when enough training data (more than 22) is presented, poly4 performs slightly better than poly4-convex. We conjecture such difference is due to the gap between sos-convex polynomials and convex polynomials (Ahmadi and Parrilo, 2012). GP has similar performance trends as poly4 but worse on average. (4) Polynomial models enjoy significant performance advantages when LS is smoother as in uniform point support (approximation of uniform patch contact). The advantage is smaller for three-point support whose LS has large flat facets.

3.3.2. Robotic experiment. We mount three screws at four different sets of locations underneath an aluminum right-angle triangular work object that weighs 1.508 kg with edge lengths of 150, 150, and 212.1 mm. The four different sets of support point locations (in millimeters) with respect to the right angle corner vertex are $[(10,10),(10,130),(130,10)]$, [(30,30), (30,90), (90,30)], [(10,10), (10,130), (90,30)], and [(30,30), (63.33,43.33), (43.33,63.33)]. Given known mass and COM projection, ideal ground truth pressure for each support point can be computed by solving three linear equations assuming each screw head approximates a point contact. Figure 4a shows a flipped view of one arrangement whose ideal LS is illustrated in Figure 4b, constructed by Minkowski addition of generalized friction at each single-point support assuming a Coulomb friction model with uniform coefficient of friction. Three pairs of symmetric facets[5] characterize indeterminate friction force when rotating about one of the three support points. Comparison among identified fourth-order homogeneous polynomials with and without convexity constraint is shown in Figure 4c and d. We can see that convex-shape constraint is essential to avoid poor generalization error when little data is available. Figure 4e and f compare the level sets of a convex quadratic (ellipsoid) and a sos-convex degree-four homogeneous polynomial, demonstrating that the higher-degree polynomial captures the facets effect better than quadratic models.

We conduct robotic poking (single-point pushing) experiments on wood and paper board surfaces. In each experiment, we generate 50 pokes (30 for training set, 10 for validation set, and 10 for test set) with randomly chosen contact points and pushing velocity directions. During each pushing action, the robot moves at a slow speed of 2.5 mm s^{-1} with a total small push-in distance of 15 mm. Each generalized velocity direction is approximated as the direction of pose displacement and generalized force is averaged over the action duration. Figure 5 shows model accuracy (averaged over four different pressure arrangements) with respect to increase in amount of training data for different methods evaluated on both the hold-out test sensor data and samples from ideal LS. We can see similar performance trends as in simulation experiments. Note that both evaluations only serve as certain reference criteria. Sensor data is noisy and all possible force measurements from a single-point pusher only cover a limited space of the set of friction loads. We also do not intend to treat the idealized LS as absolute ground truth as there is no guarantee on a uniform coefficient of friction between the support points and the underlying surface. In addition, the point contact and isotropic Coulomb friction model are only approximations of reality. Nevertheless, both evaluations demonstrate the performance advantage of our proposed poly4-cvx model.

Table 1. Comparison of average accuracy with 95% confidence interval as the amount of training data increases.

	10	20	30
poly4-cvx	**88.13**±1.80	**91.33**±1.61	**93.07**±1.45
poly4	85.27±2.12	89.40±1.98	93.00±1.62
quadratic	87.93±1.72	87.20±1.65	88.00±1.39

3.4. Stable push action generation

Prediction of the resultant object twist under a single-point push action cannot be exactly accurate. A two-point push action against an edge of the object, however, can be stable such that the object will remain attached to the pusher without slipping or breaking contact (Lynch and Mason, 1996). That is, the slider and pusher will move about the same COR point p_c. Given the level set representation $H(\mathbf{F})$, the condition for determining whether a two-point push with instantaneous generalized velocity $\mathbf{V}_{p_c}$ is stable or not is equivalent to checking whether the corresponding generalized friction load $\mathbf{F}_{p_c} = H_{inv}(\mathbf{V}_{p_c})$ lies in the applied composite wrench cone $\mathbf{F}_c$. To validate predictions based on the model, we sampled 60 random CORs and execute with the robot for three different pressure arrangements on a novel support surface material (hard poster paper). We use the same triangular block in Figure 4a with two three-point contacts $[(10,10),(10,130),(130,10)]$ and $[(30,30),(30,90),(90,30)]$ as well as full patch contact. The 60 CORs are tight rotation centers within a 400 mm × 400 mm square centered at the COM. A total of 15 out of the 60 CORs are labeled as stable. The training force–motion data are collected from pushing the object on a wood surface. Tables 1 and 2 summarize the classification accuracy and positive (stable) class recall measurements of three invertible methods with respect to increase in the amount of training data. Figure 6 shows an example (full patch contact) that the stable regions generated from the identified poly4-cvx model is much larger than the conservative analysis as in Lynch and Mason (1996), which misses the tight/closer rotation centers.

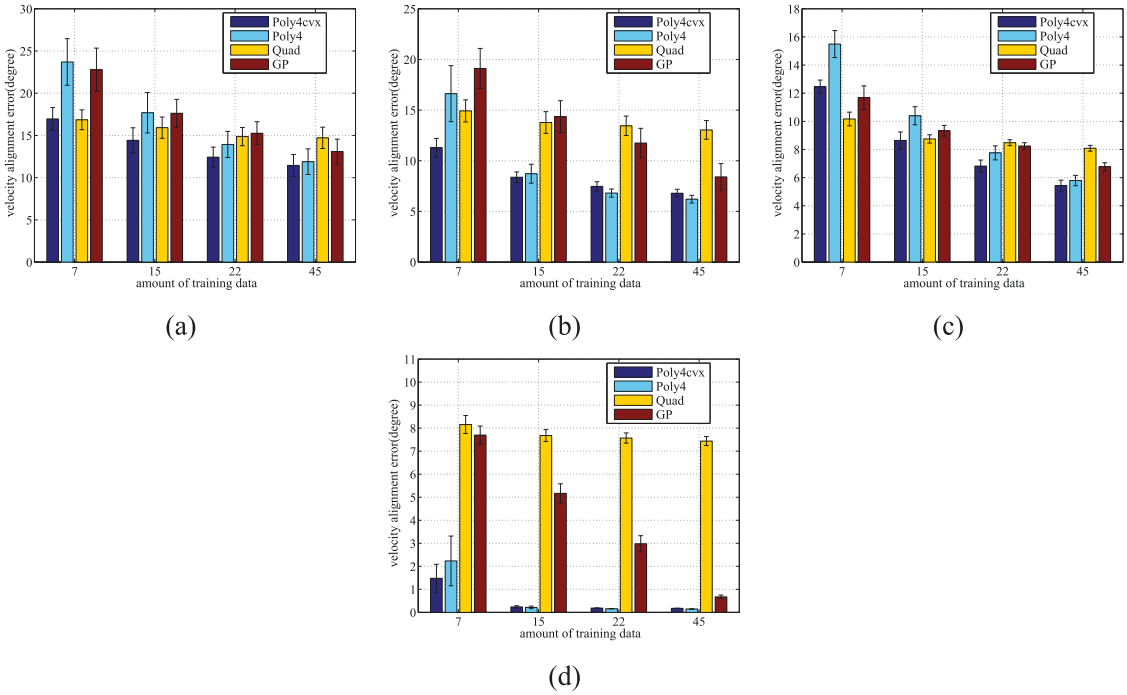

Fig. 3. Test error comparison for simulation experiments with 95% confidence bar (50 random evaluations) among different methods as amount of training data increases for three random support points and 360 support points on a ring respectively. (a) Three support points with noisy training and validation data. (b) Three support points with noise-free training and validation data. (c) Uniform circular support points with noisy training and validation data. (d) Uniform circular support points with noise-free training and validation data. Results for uniform pressure distribution within a square are similar to uniform circular support and omitted for space.

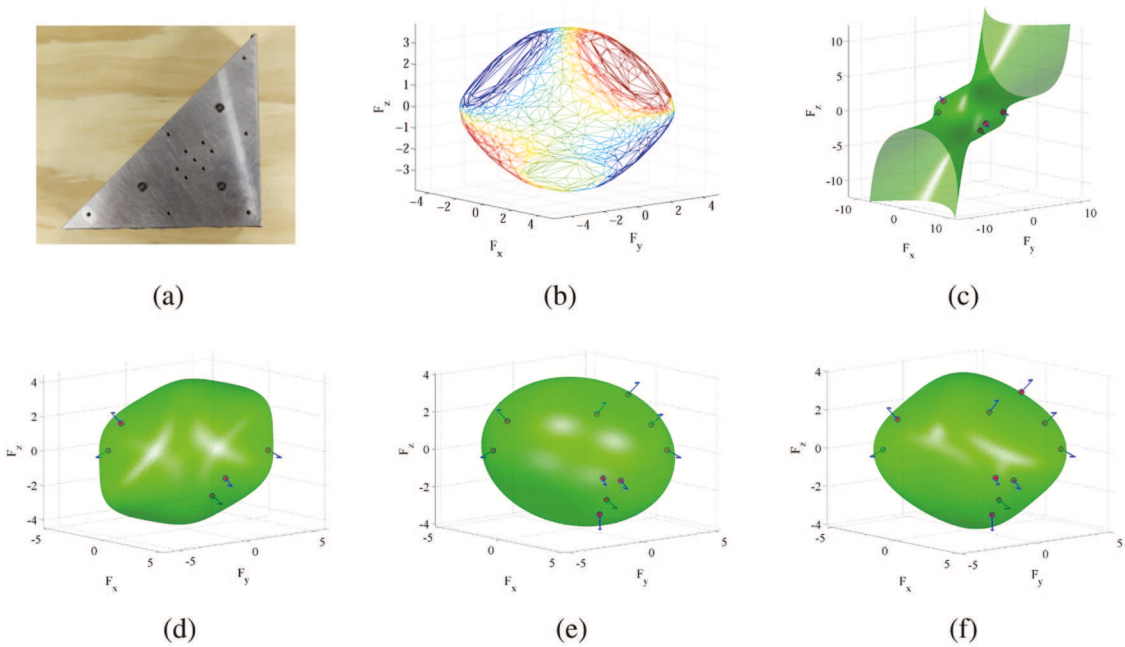

Fig. 4. Level set friction representations for the pressure arrangement in Figure 4a. (a) Triangular block with three support screws. (b) Ideal LS with facets. (c) Poly4 fit with 5 training and 5 validation data. (d) Poly4-cvx fit with 5 training and 5 validation data. (e) Convex quadratic fit with 10 training and 10 validation data. (f) Poly4-cvx fit with 10 training and 10 validation data. Red dots and blue arrows are collected generalized forces and velocities from force-torque and motion capture sensor, respectively. Parts (c), (d), (e), and (f) share the same data.

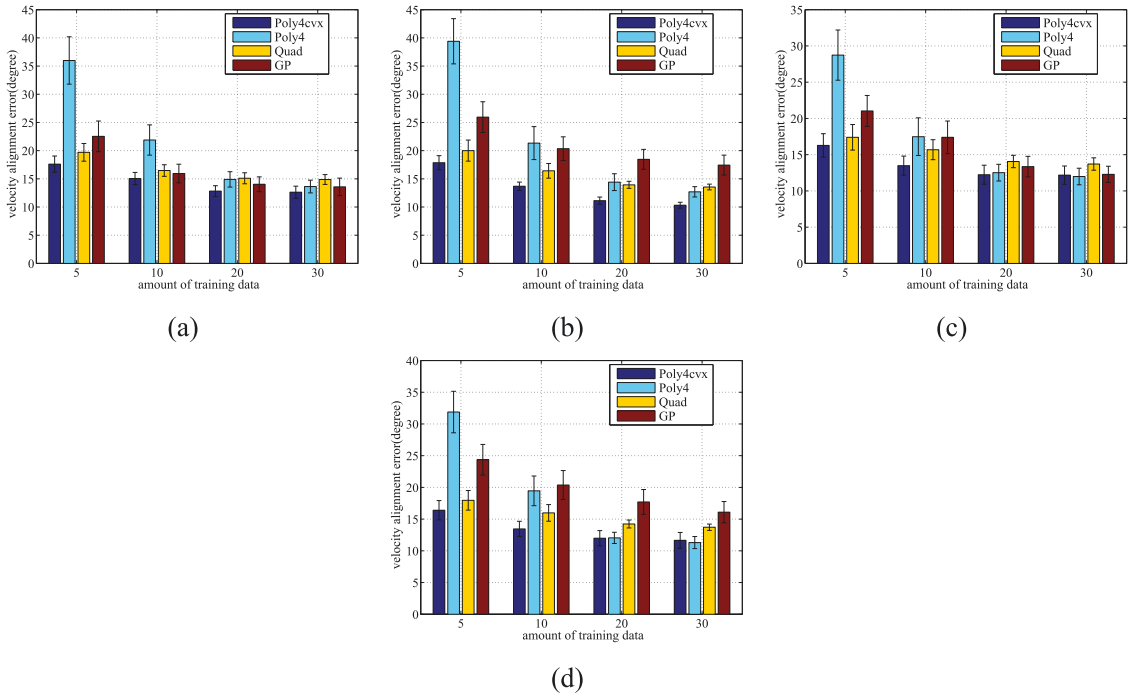

Fig. 5. Test error comparison for robotic experiments with 95% confidence bar (50 random evaluations) among different methods as the amount of training data increases for three support points (averaged over four different arrangements) on wood and hard paper board surfaces. (a) Test on sensor data (wood surface). (b) Test on data sampled from ideal LS (wood surface). (c) Test on sensor data (paper board surface). (d) Test on data sampled from ideal LS (paper board surface).

Table 2. Comparison of average positive recall with 95% confidence interval as the amount of training data increases.

	10	20	30
poly4-cvx	**90.13**±3.54	**96.69**±1.93	**98.18**±1.32
poly4	79.96±5.25	92.76±2.90	97.18±1.84
quadratic	73.18±4.61	73.38±4.69	73.87±4.63

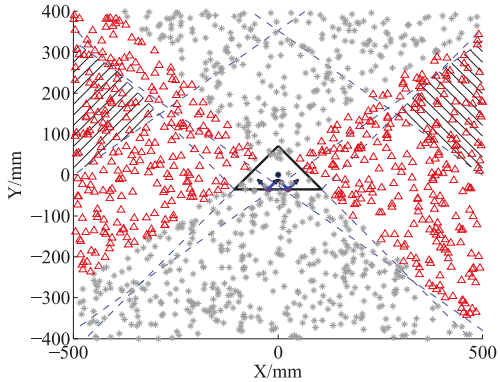

Fig. 6. Hatched areas correspond to stable COR regions based on the conservative analysis (Lynch and Mason, 1996). Red triangles are stable CORs and gray stars are non-stable CORs based on the poly4-cvx model. The two push points are 50 mm wide. The pusher and the object are covered with electrical tape and gaffer tape, respectively, with measured coefficient of friction equaling one.

4. Kinematic contact modeling

With a position-controlled manipulator, we are given a single-point finger contact at $\mathbf{p}$ with inward normal $\mathbf{n}_p$, pushing velocity $\mathbf{v_p}$ and coefficient of friction μ_c between the pusher and the object. The task is to resolve the incurred body twist $\mathbf{V}$ and contact mode (sticking, slipping, breaking contact): find a $\mathbf{V}$ consistent with the contact mode at $\mathbf{p}$ while the applied wrench, solved as an intermediate output (not a supplied control), equals the corresponding generalized friction load.

4.1. Single-point pusher

We introduce the concept of a motion cone first proposed in Mason (1986a). Let the Jacobian matrix $J_p = \begin{bmatrix} 1 & 0 & -p_y \\ 0 & 1 & p_x \end{bmatrix}$, and denote by $\mathbf{F}_l = J_p^{\mathrm{T}}\mathbf{f}_l$ and $\mathbf{F}_r = J_p^{\mathrm{T}}\mathbf{f}_r$ the left and right edges of the applied wrench cone with corresponding resultant twist directions $\mathbf{V}_l = \nabla H(\mathbf{F}_l)$ and $\mathbf{V}_r = \nabla H(\mathbf{F}_r)$, respectively. The left edge of the motion cone is $\mathbf{v}_l = J_p\mathbf{V}_l$ and the right edge of the motion cone is $\mathbf{v}_r = J_p\mathbf{V}_r$. Mason (1986a) showed that if the contact point pushing velocity $\mathbf{v_p}$ is inside the motion cone, i.e. $\mathbf{v_p} \in \mathbf{K}(\mathbf{v}_l, \mathbf{v}_r)$, the contact sticks. When $\mathbf{v_p}$ is outside the motion cone, sliding occurs. If $\mathbf{v_p}$ is to the left of $\mathbf{v}_l$, the pusher will slide left with respect to the object. Otherwise,

if $\mathbf{v_p}$ is to the right of $\mathbf{v}_r$, the pusher will slide right as shown in Figure 7.

The following constraints hold assuming sticking contact:

$$v_{px} = V_x - \omega p_y \tag{9}$$

$$v_{py} = V_y + \omega p_x \tag{10}$$

$$\mathbf{V} = \frac{1}{2} k \cdot \nabla H(\mathbf{F}), \quad k > 0 \tag{11}$$

$$\tau = -p_y F_x + p_x F_y \tag{12}$$

In the case of ellipsoid (convex quadratic) representation, i.e. $H(\mathbf{F}) = \mathbf{F}^{\mathrm{T}} A \mathbf{F}$ where A is a positive-definite matrix, the problem is a full rank linear system with a unique solution. Lynch et al. (1992) give an analytical solution when A is diagonal. We show that a unique analytic solution exists for any positive-definite symmetric matrix A. Let $\mathbf{t} = [-p_y, p_x, -1]^{\mathrm{T}}$, Equations (9)–(12) can be written as

$$J_p \mathbf{V} = \mathbf{v_p} \tag{13}$$

$$\mathbf{V} = kA\mathbf{F} \tag{14}$$

$$\mathbf{t}^T F = 0 \tag{15}$$

Using the notation $D = [J_p^{\mathrm{T}}, A^{-1}\mathbf{t}]^{\mathrm{T}}$ and $\mathbf{V}_p = [\mathbf{v_p}^T, 0]^{\mathrm{T}}$, we can combine the above equations into one linear equation:

$$\mathbf{V} = D^{-1} \mathbf{V}_p \tag{16}$$

Theorem 3. *Pushing with single sticking contact and the convex quadratic representation of LS (abbreviated as P.1) has a unique solution from a linear system.*

Proof. From Equation (16), we only need to prove that D is invertible. ■

1. The row vectors of J_p are linearly independent and span a plane.
2. Here $J_p \mathbf{t} = 0$ implies $\mathbf{t}$ is orthogonal to the spanned plane.
3. If D is not full rank, then $A^{-1}\mathbf{t}$ must lie in the spanned plane and is therefore orthogonal to $\mathbf{t}$. This contradicts the fact that $\langle \mathbf{t}, A^{-1}\mathbf{t} \rangle > 0$ for positive-definite matrix A^{-1} and non-zero vector $\mathbf{t}$.

Corollary 1. *Pushing with single sticking contact and the general homogeneous convex polynomial representation of LS is reducible to solving a sequence of subproblems P.1.*

For general convex polynomial representation $H(\mathbf{F})$, the following optimization is equivalent to Equations (9)–(12):

$$\underset{\mathbf{F}}{\text{minimize}} \qquad \|J_p \nabla H(\mathbf{F}) - \mathbf{v_p}\| \tag{17}$$

$$\text{subject to} \qquad \mathbf{t}^{\mathrm{T}} F = 0 \tag{18}$$

When $H(\mathbf{F})$ is of the convex quadratic (ellipsoidal) form, the analytical minimizer is $\mathbf{F} = A^{-1} D^{-1} \mathbf{V}_p$. In the case of high-order convex homogeneous polynomials, we can resort to an iterative solution where we use the Hessian matrix as a local ellipsoidal approximation, i.e. set $A_t = \nabla^2 H(\mathbf{F_t})$ and compute $\mathbf{F_{t+1}} = A_t^{-1} D^{-1} \mathbf{V}_p$ until convergence.

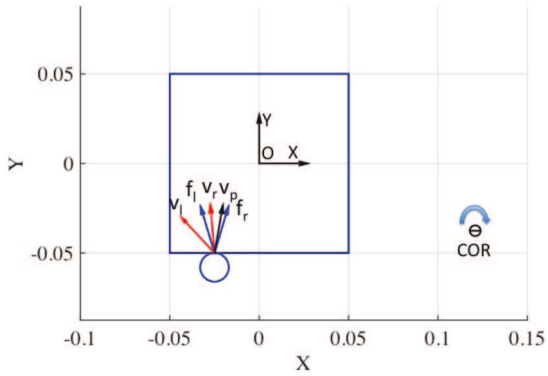

Fig. 7. Mechanics of single-point pushing. The square has a uniform pressure distribution over 100 support grid points sharing the same coefficient of friction. We use a fourth-order convex polynomial to represent the LS. The finger's pushing velocity is to the right of the motion cone and, hence, the finger will slide to the right. The instantaneous COR, computed using the model described in Section 4.1, is marked as a circle with a negative sign indicating clockwise rotation.

When $\mathbf{v_p}$ is outside of the motion cone, assuming right sliding occurs without loss of generality, the wrench applied by the finger equals $\mathbf{F}_r$. The resultant object twist $\mathbf{V}$ follows the same direction as $\mathbf{V}_r$ with proper magnitude such that the contact is maintained:

$$\mathbf{V} = s\mathbf{V}_r \tag{19}$$

$$s = \frac{\mathbf{n_p}^{\mathrm{T}} \mathbf{v}_p}{\mathbf{n_p}^{\mathrm{T}} \mathbf{v}_l} \tag{20}$$

4.2. Multi-contacts

Mode enumeration is tedious for multiple contacts. The linear complementarity formulation for frictional contacts (Stewart and Trinkle, 1996) provides a convenient representation. Denote by m the total number of contacts, the quasi-static force–motion equation is given by

$$\mathbf{V} = k \nabla H(\mathbf{F}) \tag{21}$$

where the total applied wrench is the sum of normal and frictional wrenches over all applied contacts:

$$\mathbf{F} = \sum_{i=1}^{m} J_{p_i}^{\mathrm{T}} (f_{n_i} \mathbf{n_{p_i}} + D_{\mathbf{p_i}} \mathbf{f}_{t_i}) \tag{22}$$

Here f_{n_i} is the normal force magnitude along the normal $\mathbf{n_i}$, and $\mathbf{f}_{t_i}$ is the vector of tangential friction force magnitudes along the column vector basis of $D_{\mathbf{p_i}} = [\mathbf{t}_{p_i}, -\mathbf{t}_{p_i}]$. The velocity at contact point $\mathbf{p_i}$ on the object is given by $J_{p_i} \mathbf{V}$.

The first-order complementarity constraints on the normal force magnitude and the relative velocity are given by

$$0 \leq f_{n_i} \perp (\mathbf{n}_{p_i}^{\mathrm{T}}(J_{p_i}\mathbf{V} - \mathbf{v}_p)) \geq 0 \tag{23}$$

The complementarity constraints for Coulomb friction are given by

$$0 \leq \mathbf{f}_{t_i} \perp (D_{\mathbf{p_i}}^{\mathrm{T}}(J_{p_i}\mathbf{V} - \mathbf{v}_p) + \mathbf{e}\lambda_i) \geq 0 \tag{24}$$

$$0 \leq \lambda_i \perp (\mu_i f_{n_i} - \mathbf{e}^{\mathrm{T}}\mathbf{f}_{t_i}) \geq 0 \tag{25}$$

where μ_i is the coefficient of friction at $\mathbf{p_i}$ and $\mathbf{e} = [1; 1]$. In the case of ellipsoid (convex quadratic) representation, i.e. $H(\mathbf{F}) = \mathbf{F}^{\mathrm{T}}A\mathbf{F}$ where A is a positive-definite matrix, Equations (21)–(25) can be written in matrix form:

$$\begin{bmatrix} 0 \\ \alpha \\ \beta \\ \gamma \end{bmatrix} = \begin{bmatrix} A^{-1}/k & -N^{\mathrm{T}} & -L^{\mathrm{T}} & 0 \\ N & 0 & 0 & 0 \\ L & 0 & 0 & E \\ 0 & \mu & -E^{\mathrm{T}} & 0 \end{bmatrix} \begin{bmatrix} \mathbf{V} \\ \mathbf{f}_n \\ \mathbf{f}_t \\ \lambda \end{bmatrix} + \begin{bmatrix} 0 \\ \mathbf{a} \\ \mathbf{b} \\ 0 \end{bmatrix} \tag{26}$$

$$0 \leq \begin{bmatrix} \alpha \\ \beta \\ \gamma \end{bmatrix} \perp \begin{bmatrix} \mathbf{f}_n \\ \mathbf{f}_t \\ \lambda \end{bmatrix} \geq 0$$

where the binary matrix $E \in R^{2m \times m}$ equals $\begin{bmatrix} \mathbf{e} & & \\ & \ddots & \\ & & \mathbf{e} \end{bmatrix}$, $\mu = [\mu_1, \ldots, \mu_m]^{\mathrm{T}}$, the stacking matrix $N \in R^{m \times 3}$ equals $[\mathbf{n}_{p_1}^{\mathrm{T}}J_{p_1}; \ldots; \mathbf{n}_{p_m}^{\mathrm{T}}J_{p_m}]$, the stacking matrix $L \in R^{2m \times 3}$ equals $[D_{p_1}^{\mathrm{T}}J_{p_1}; \ldots; D_{p_m}^{\mathrm{T}}J_{p_m}]$, the stacking vector $\mathbf{a} \in R^m$ equals $[-\mathbf{n}_{p_1}^{\mathrm{T}}\mathbf{v}_{p_1}, \ldots, -\mathbf{n}_{p_m}^{\mathrm{T}}\mathbf{v}_{p_m}]^{\mathrm{T}}$, and the vector $\mathbf{b} \in R^{2m}$ equals $[-D_{p_1}^{\mathrm{T}}\mathbf{v}_{p_1}, \ldots, -D_{p_m}^{\mathrm{T}}\mathbf{v}_{p_m}]^{\mathrm{T}}$.

Note that the positive scalar k will not affect the solution value of $\mathbf{V}$ since $\mathbf{f}_n$ and $\mathbf{f}_t$ will scale accordingly. Hence, we can drop the scalar k and further substitute $\mathbf{V} = A(N^{\mathrm{T}}\mathbf{f}_n + L^{\mathrm{T}}\mathbf{f}_t)$ into Equation (26) and reach the standard linear complementarity form as follows:

$$\begin{bmatrix} \alpha \\ \beta \\ \gamma \end{bmatrix} = \begin{bmatrix} NAN^{\mathrm{T}} & NAL^{\mathrm{T}} & 0 \\ LAN^{\mathrm{T}} & LAL^{\mathrm{T}} & E \\ \mu & -E^{\mathrm{T}} & 0 \end{bmatrix} \begin{bmatrix} \mathbf{f}_n \\ \mathbf{f}_t \\ \lambda \end{bmatrix} + \begin{bmatrix} \mathbf{a} \\ \mathbf{b} \\ 0 \end{bmatrix} \tag{27}$$

$$0 \leq \begin{bmatrix} \alpha \\ \beta \\ \gamma \end{bmatrix} \perp \begin{bmatrix} \mathbf{f}_n \\ \mathbf{f}_t \\ \lambda \end{bmatrix} \geq 0$$

Similarly, for the case of high-order convex homogeneous polynomials, we can iterate between taking the linear Hessian approximation and solving the LCP problem in Equation (27) until convergence.

Lemma 1. *The object is quasi-statically jammed or grasped if Equation* (27) *yields no solution.*

Figure 8 provides a graphical proof. When Equation (27) yields no solution, either there is no feasible kinematic motion of the object without penetration or all the friction loads associated with the feasible instantaneous twists cannot balance against any element from the set of possible applied wrenches. In this case, the object is quasi-statically jammed or grasped between the fingers. Neither the object nor the end effector can move.

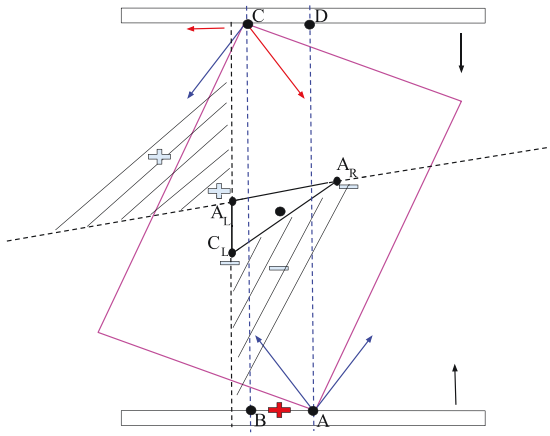

Fig. 8. Using moment labeling (Mason, 2001), the COR has positive sign (counter-clockwise) and can only lie in the band between the two blue contact normal lines. Further, the COR must lie on segment AB (contact point A sticks) or segment CD (contact point C sticks) since otherwise both contacts will slip, but the total wrench from the two left edges of the friction cones has negative moment that cannot cause counter-clockwise rotation. Without loss of generality, we can assume COR (red plus) lies on segment AB, leading to sticking contact at A and left sliding at C. Following a similar analysis using the force dual graphical approach (Brost and Mason, 1991), each single friction force can be mapped to its instantaneous resultant signed COR whose convex combination forms the set of all possible CORs under the composite friction forces. The COR can either be of positive sign in the upper left hatched region or negative sign in the lower right hatched region that contradicts the proposed AB segment. Hence, jamming occurs and neither the gripper nor the object can move. This corresponds exactly to the no solution case of Equation (27).

4.3. Stochastic Modeling

Frictional interaction is inherently stochastic. Two major sources contribute to the uncertainty in planar motion: (1) indeterminacy of the supporting friction distribution $\mathbf{f_r}$ due to changing pressure distribution and coefficients of friction between the object and support surface; (2) the coefficient of friction μ_c between the object and the robot end effector. We sample μ_c uniformly from a given range and model the effect of changing support friction distribution by sampling the parameters a in $H(\mathbf{F}; a)$ from a distribution that satisfies the following.

1. Samples from the distribution should result in an even-degree homogeneous convex polynomial to represent the LS.

Table 3. Average deviation metric (in millimeters) between the simulated final pose and actual final pose with 95% confidence interval. The third, sixth, and ninth rows are the deviation from the ground truth initial pose and final pose to indicate how much the object is moved due to the push. In most cases, the fourth-order convex (poly4) polynomial has better accuracy. The average normalized percentage error for poly4 is 20.05% and for quadratic is 21.39%. However, the accuracy of a fixed deterministic model is bounded by the inherent variance of the system.

	rect1	rect2	rect3	tri1	tri3	ellip1	ellip2	ellip3	hex	butter
poly4-delrin	8.28±0.29	5.37±0.23	6.10±0.21	9.71±0.33	7.54±0.23	7.68±0.51	8.90±1.40	7.35±0.38	6.38±0.28	4.83±0.27
quad-delrin	8.60±0.35	5.92±0.14	8.20±0.16	9.90±0.41	8.18±0.15	6.85±0.25	6.29±0.24	8.08±0.51	6.42±0.12	5.97±0.23
delrin	35.48	40.53	35.98	36.91	34.66	32.18	38.05	33.37	33.55	34.09
poly4-abs	5.86±0.11	7.48±0.80	3.59±0.12	7.13±0.26	5.17±0.38	8.45±1.13	9.18±1.26	5.93±0.19	7.56±0.39	3.94±0.11
quad-abs	6.07±0.16	6.74±0.27	6.19±0.18	8.00±0.37	7.17±0.37	6.66±0.28	7.69±0.27	5.78±0.21	8.19±0.21	5.39±0.15
abs	34.14	39.74	33.98	35.43	32.37	32.68	33.53	32.45	33.23	33.53
poly4-plywood	6.86±0.71	6.86±0.13	5.93±0.33	4.61±0.13	7.21±0.47	4.39±0.16	4.99±0.31	5.72±0.31	8.41±0.24	4.72±0.17
quad-plywood	6.20±0.20	7.22±0.18	6.88±0.18	5.96±0.19	9.43±0.56	4.42±0.12	5.84±0.20	6.46±0.26	8.85 ±0.17	6.05±0.22
plywood	31.86	33.22	32.94	32.81	33.78	27.24	28.23	33.29	32.77	34.10

2. The mean can be set as a prior estimate and the amount of variance controlled by one parameter.

The n_{df} degree of freedom Wishart distribution $S \sim W(\hat{S}, n_{df})$ (Wishart, 1928) with mean $n_{df}S_{est}$ and variance $Var(S_{ij}) = n_{df}(\hat{S}_{ij}^2 + \hat{S}_{ii}\hat{S}_{jj})$ is defined over symmetric positive-semidefinite random matrices as a generalization of the chi-squared distribution to multi-dimensions. For ellipsoidal (convex quadratic) $H(\mathbf{F}; A)$, we can directly sample from $\frac{1}{n_{df}}W(A_{est}, n)$ where A_{est} is some estimated value from data or fitted for a particular pressure distribution. Sampling from general convex polynomials is hard. Fortunately, we find that sampling from the sos-convex (Magnani et al., 2005; Parrilo, 2000) polynomials subset is not. The key is the coefficient vector a of a sos-convex polynomial $H(\mathbf{F}; a)$ has a unique one-to-one mapping to a positive-definite matrix Q so that we can first sample $\tilde{Q}$ from[6] $\frac{1}{n_{df}}W(Q, n_{df})$ and then map back to $\tilde{a}$ through Equation (5). The degree of freedom parameter n_{df} determines the sampling variance. The smaller n_{df} is, the noisier the system will be.

4.4. Deterministic pushing model evaluation

We evaluate the single-contact deterministic model on the large-scale MIT pushing dataset (Yu et al., 2016) and the data from the identification experiments in Section 3.3.2. For the MIT pushing dataset, we use 10 mm s^{-1} velocity data logs for 10 objects[7] on the hard surfaces including delrin, abs, and plywood. The force torque signal is first filtered with a low-pass filter and 5 wrench-twist pairs evenly spaced in time are extracted from each push action log file. A total of 10 random train–test splits (20% of the logs for training, 10% for validation, and the remainder for testing) are conducted for each object–surface scenario.

Given two poses $q_1 = [x_1, y_1, \theta_1]$ and $q_2 = [x_2, y_2, \theta_2]$, we define the deviation metric $d(q_1, q_2)$ which combines both the displacement and angular offset as $d(q_1, q_2) = \sqrt{(x_1 - x_2)^2 + (y_1 - y_2)^2} + \rho \cdot \min(|\theta_1 - \theta_2|, 2\pi - |\theta_1 - \theta_2|)$, where ρ is the characteristic length of the object (e.g. radius of gyration or radius of minimum circumscribed circle). A one-dimensional coarse grid search over the coefficient of friction μ_c between the pusher and object is chosen to minimize average deviation of the predicted final pose and ground truth final pose on training data. Table 3 shows the average metric with a 95% confidence interval. Interestingly, we find that using more training data does not improve the performance much. This is likely due to the inherent stochasticity and changing surface conditions as reported in Yu et al. (2016).

The well-machined objects in the MIT pushing dataset are close to uniform patch pressure. We also test on discrete pressures. The triangular object used in the identification experiments in Section 3.3.2 are given different configurations of three-point discrete support. We use wrench twists pairs sampled from the ideal LS for training. The coefficient of friction between the object and pusher is determined by a grid search over 40% of the logs. We use the remaining 60% to evaluate simulation accuracy. Results are reported in Table 4.

Table 4. Average deviation (in millimeters) between the simulated final pose and actual final pose with 95% confidence interval for three-point support. The wrench–twist pairs used for training the model are generated from the ideal LS. The third and sixth rows are the deviation from the ground truth initial pose and final pose to indicate how much the object is moved due to the push. The fourth-order convex (poly4) polynomial has better accuracy for each pressure–surface combination. The average normalized error for poly4 is 20.48% and for quadratic is 24.97%.

	3pts1	3pts2	3pts3	3pts4
poly4-hardboard	3.52±0.21	2.75±0.25	2.92±0.27	2.80±0.23
quad-hardboard	3.82±0.24	3.63±0.27	3.35±0.23	3.96±0.28
hardboard	16.63	13.86	14.83	15.15
poly4-plywood	3.78±0.11	2.80±0.15	2.84±0.16	3.26±0.11
quad-plywood	4.24±0.15	3.56±0.17	3.28±0.08	4.12±0.13
plywood	16.56	13.81	15.27	14.20

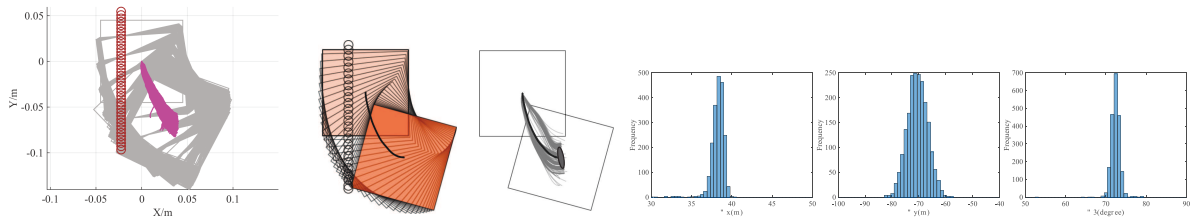

Fig. 9. Stochastic modeling of single-point pushing with the fourth-order sos-convex polynomial representation of the LS using wrench–twist pairs generated from 64 grids with uniform pressure. (a) Stochastic simulation results. (b) Figure 9 of Yu et al. (2016), reprinted with permission. (c) Δx histogram. (d) Δy histogram. (e) $\Delta\theta$ histogram. The degree of freedom in the sampling distribution equals 20. The contact coefficient of friction between the pusher and the object is uniformly sampled from 0.15 to 0.35. The trajectories are qualitatively similar to the experimental results in Figure 9 of Yu et al. (2016).

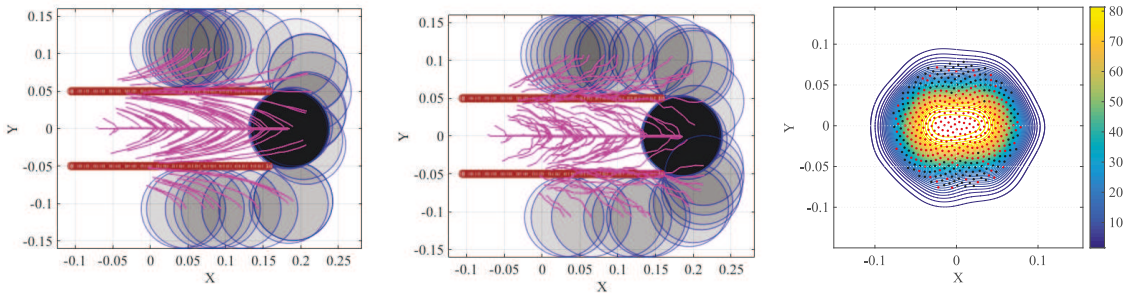

Fig. 10. Simulation results using the proposed contact model illustrating the process of two-point fingers pushing a circle to reduce initial pose uncertainty. A total of 500 initial object center positions are uniformly sampled from a circle of radius 7.88 cm. (a) A total of 100 pushed trajectories of different initial poses using ellipsoid representation of $H(\mathbf{F})$ with $n_{df} = 200$. (b) A total of 100 pushed trajectories of different initial poses using ellipsoid representation of $H(\mathbf{F})$ with $n_{df} = 10$ (larger noise). (c) Kernel density plot of the convergence region for $n_{df} = 10$. Convergent initial poses are in red and the rest are in black.

4.5. Stochastic pushing model simulation

Yu et al. (2016) reported the same 2000 straight-line pushes in a highly controlled setting result in a distribution of final poses, demonstrating the inherent stochastic nature of pushing. We perform simulations using the same object and pusher geometry and push distance. The 2000 resultant trajectories and histogram plot of pose changes are shown in Figure 9. We note that although the mean and variance pose changes are similar to the experiments with abs material in Yu et al. (2016), the distribution resembles a single Gaussian distribution which differs from the multiple-mode distribution in Figure 10 of Yu et al. (2016). We conjecture this is due to a time-varying stochastic process where coefficients of friction between surfaces drift due to wear.

Using a high-fidelity contact model is important to evaluate and generate plans for uncertainty reduction. We also simulate the effects of initial pose uncertainty reduction with two-point fingers under the stochastic contact model. The circular object has a radius of 5.25 cm. The two fingers separated by 10 cm perform a straight-line push of 26.25 cm. The desired goal is to have the object centered with respect to the two fingers. Figure 10a and b compare the resultant trajectories under different amount of system noise. We find that despite larger noise in the resultant trajectories, the convergent region of the stable goal pose differs by less than 5% and the difference is mostly around the uncertainty boundary. A kernel density plot of the convergence region is shown in Figure 10c for $n_{df} = 10$. This demonstrates multiple constraints can induce a large region of attraction despite uncertainty.

4.6. Grasping under uncertainty

We conducted robotic experiments to evaluate our contact model for grasping. Figure 11a shows two rectangular objects with the same geometry but different pressure distributions. Another experiment is conducted for a butterfly-shaped object shown in Figure 12a. We use the Robotiq C-85 2-finger gripper (ROBOTIQ, 2017) and represent it as a planar parallel-jaw gripper with rectangular fingers in simulation as shown in Figure 11e and e. Convex quadratic representations of $H(\mathbf{F})$ are trained from wrench–twist pairs assuming a uniform friction distribution along the object boundary. The sampling degree of freedom n_{df} equals 250 with contact friction coefficient μ_c sampled uniformly from

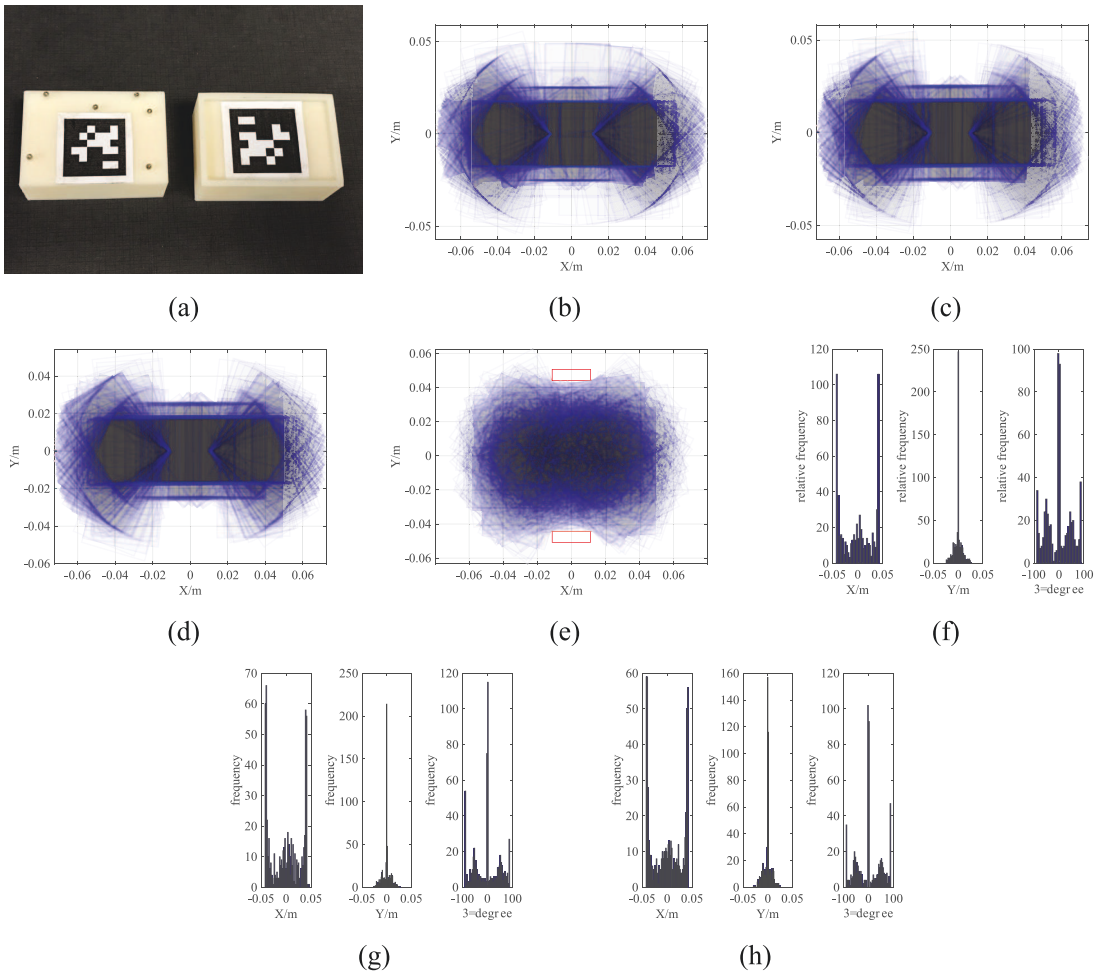

Fig. 11. Experiments on the rectangular objects with different pressure distributions. A total of 600 initial poses are sampled whose centers are uniformly distributed in a circle of radius of 20 mm and angles are uniformly distributed from -90° to 90°. (a) Two 50 mm $\times$ 35 mm rectangles with six points and boundary pressure distribution. (b) Distribution of the simulated post-grasp poses using the stochastic contact model. (c) Distribution of the experimental post-grasp poses for the boundary pressure. (d) Distribution of the experimental post-grasp poses for the six-point pressure. (e) Initial uncertainty of 600 poses. (f) Histogram of the simulated post-distribution. (g) Histogram of the experimental post-distribution for the boundary pressure. (h) Histogram of the experimental post-distribution for the six-point pressure.

$[0.015, 0.02]$. For the rectangles of both pressure distributions, the simulated results with the stochastic contact model match well with the experimental data. Although exact matching of the distribution is hard to achieve, the different modes corresponding to different final post-grasp contact conditions can be well captured to plan a sequence of grasps that shrink the uncertainty to a singleton (Zhou et al., 2017). However, the model fails to capture the stability of grasps and the deformation of objects. In the case of a butterfly-shaped object, many unstable grasps and jamming equilibria exist, but as the fingers increase the gripping force the object will "fly" away as the stored elastic energy turns into large accelerations that violates the quasi-static assumptions of our model, as revealed in the scattered post-grasp distribution in Figure 12c. We also compare the cases where dynamics do not play a major role: Figure 12d shows the zoomed in plots to compare with simulation results in Figure 12b. Despite qualitative similarity, the simulation results deviate more compared with the case for the rectangular geometry. As shown in the histogram plots in Figure 12f and h, the simulation returns more jamming and grasping final states as illustrated by the spikes in θ.

5. Conclusions and future work

In this paper, we have proposed the use of sub-level sets and gradients of a function to represent rigid-body planar friction loads and velocities, respectively. The maximum work inequality implies that such a function needs to be convex. We additionally require the properties of symmetry, scale invariance, and efficient invertibility that lead us to choose a convex even-degree homogeneous polynomial

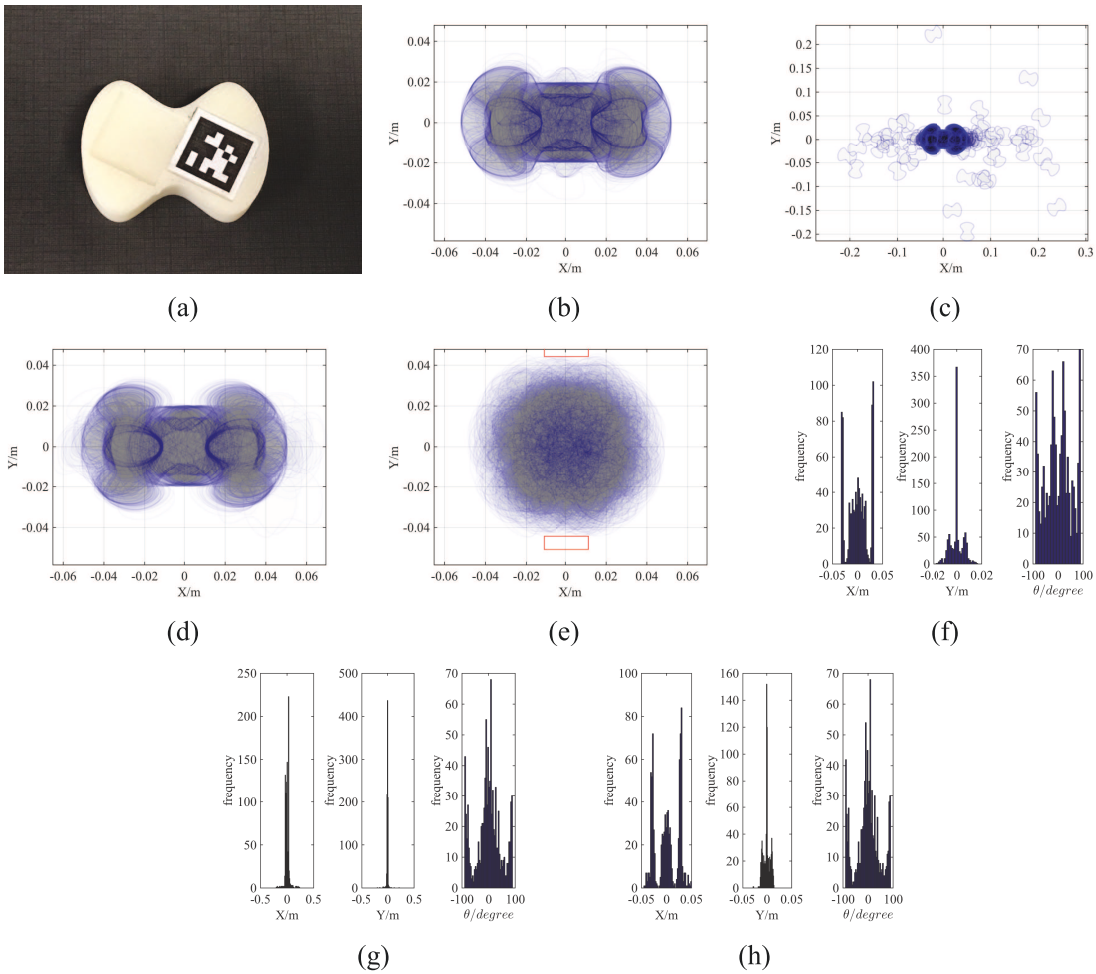

Fig. 12. Experiments on the butterfly object. The longer diameter between the convex curves is 39 mm and the shorter diameter between the concave curves is 28.6 mm. A total of 900 initial poses are sampled where the centers lie uniformly in a circle of radius 30 mm and the frame angles are uniformly distributed in -90° to 90°. (a) Butterfly-shaped object with boundary pressure distribution used for the experiment. (b) Distribution of the simulated post-grasp poses using the stochastic contact model. (c) Distribution of the object poses after the grasping actions from experimental data. (d) Zoomed-in distribution of the object poses after the grasping actions around the origin. (e) Initial uncertainty of 900 poses. (f) Histogram plot for the simulated post-distribution. (g) Histogram plot of the experimental post-distribution. (h) Histogram plot of the experimental post-distribution around the origin.

representation. The model enjoys the benefits of both accuracy and data efficiency. We then extended the model to the kinematic contact level that resolves the object twist and contact mode given the velocity input (single and multiple) from a position-controlled manipulator. To model the inherent uncertainty in frictional mechanics, we derive methods that enable sampling from the family of sos-convex polynomials. The models and applications have been validated with large-scale robotic pushing and grasping experiments.

Although physics-inspired models are data-efficient and easily generalizable, we also see the limitation of the imposing assumptions, e.g. failures in the butterfly-shaped object-grasping experiment due to the first-order quasi-static model. Much work remains to be done. On the simulator end: (1) how to increase the accuracy without losing convergence speed for a high-order polynomial-based representation of $H(\mathbf{F})$; and (2) how to handle penetration properly when the integration step is large. On the application side: (1) how to quickly identify both the mean and variance of the sampling distribution to match with experimental data; and (2) how to plan a robust sequence of grasp and push actions for uncertainty reduction using the stochastic contact model. In the long term, we are interested in combining learning-based and physics-based techniques to synthesize robust strategies for contact-rich manipulation through exploiting task mechanics.

Acknowledgements

The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Army Research Laboratory of the US Government. The US Government is authorized to reproduce and distribute reprints for Government

purposes notwithstanding any copyright notation herein. Discussions with Alberto Rodriguez and Kuan-Ting Yu motivated the stochastic simulation of pushing.

Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This work was conducted in part through collaborative participation in the Robotics Consortium sponsored by the US Army Research Laboratory under the Collaborative Technology Alliance Program (Cooperative Agreement W911NF-10-2-0016) and National Science Foundation (grant number IIS-1409003).

Notes

1. Our open-source simulation software and data are available at: https://github.com/robinzhoucmu/Pushing
2. Throughout the paper, we use a local coordinate frame with the origin set as the projection of the COM onto the supporting surface. However, the choice of the origin can be any other point of convenience.
3. In practice, $\mathbf{F}_t$ will not be near the point of origin whose $H(\mathbf{F}_t)$ is all zero. A small diagonal regularization can be added to $H(\mathbf{F}_t)$ before inversion to improve numerical stability.
4. The squared exponential kernel gives better performance over linear and polynomial. Normalizing the input load to a unit vector improves performance by requiring the GP to ignore scale. Every $(\mathbf{F}, \mathbf{V})$ input pair is augmented with $(-\mathbf{F}, -\mathbf{V})$ for training.
5. The third one is in the back not visible from presented view.
6. We note that adding a small constant on the diagonal elements of $\tilde{Q}$ improves numerical stability.
7. Despite having the same experimental set up and similar geometry and friction property to the other two triangular shapes, the results for object Tr2 is about 1.5–2 times worse. Owing to time constraints, we have not ruled out the possibility that the data for object Tr2 is corrupted.

References

Ahmadi AA and Parrilo PA (2012) A convex polynomial that is not sos-convex. *Mathematical Programming* 135: 275–292.

Akella S and Mason MT (1998) Posing polygonal objects in the plane by pushing. *The International Journal of Robotics Research* 17(1): 70–88.

Boyd S and Vandenberghe L (2004) *Convex Optimization*. Cambridge: Cambridge University Press.

Brost RC (1988) Automatic grasp planning in the presence of uncertainty. *The International Journal of Robotics Research* 7(1): 3–17.

Brost RC and Mason MT (1991) Graphical analysis of planar rigid-body dynamics with multiple frictional contacts. In: *The fifth international symposium on Robotics research*. Cambridge, MA: MIT Press, pp. 293–300.

Cole A, Hsu P and Sastry S (1992) Dynamic control of sliding by robot hands for regrasping. In: *IEEE Transactions on Robotics and Automation* 8(1): 42–52.

Dogar M and Srinivasa S (2010) Push-grasping with dexterous hands: Mechanics and a method. In: *Proceedings of 2010 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2010)*.

Erdmann M (1986) Using backprojections for fine motion planning with uncertainty. *The International Journal of Robotics Research* 5(1): 19.

Erdmann M (1994) On a representation of friction in configuration space. *The International Journal of Robotics Research* 13(3): 240–271.

Ferrari C and Canny J (1992) Planning optimal grasps. In: *Proceedings 1992 IEEE International Conference on Robotics and Automation (1992)*. IEEE, pp. 2290–2295.

Goyal S, Ruina A and Papadopoulos J (1991) Planar sliding with dry friction. Part 1. Limit surface and moment function. *Wear* 143: 307–330.

Hogan FR and Rodriguez A (2016) Feedback control of the pusher–slider system: A story of hybrid and underactuated contact dynamics. *arXiv preprint arXiv:1611.08268*.

Howe RD and Cutkosky MR (1996) Practical force-motion models for sliding manipulation. *The International Journal of Robotics Research* 15(6): 557–572.

Kopicki M, Zurek S, Stolkin R, Mörwald T and Wyatt J (2011) Learning to predict how rigid objects behave under simple manipulation. In: *2011 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, pp. 5722–5729.

Lozano-Perez T, Mason MT and Taylor RH (1984) Automatic synthesis of fine-motion strategies for robots. *The International Journal of Robotics Research* 3(1): 3–24.

Lynch K, Maekawa H and Tanie K (1992) Manipulation and active sensing by pushing using tactile feedback. In: *Proceedings of the 1992 lEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 1. IEEE, pp. 416–421.

Lynch KM (1993) Estimating the friction parameters of pushed objects. In: *Proceedings of the 1993 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS'93)*, volume 1. IEEE, pp. 186–193.

Lynch KM and Mason MT (1996) Stable pushing: Mechanics, controllability, and planning. *The International Journal of Robotics Research* 15(6): 533–556.

Magnani A, Lall S and Boyd S (2005) Tractable fitting with convex polynomials via sum-of-squares. In: *44th IEEE Conference on Decision and Control, 2005 and 2005 European Control Conference (CDC-ECC'05)*. IEEE, pp. 1672–1677.

Mason MT (1986a) Mechanics and planning of manipulator pushing operations. *The International Journal of Robotics Research* 5(3): 53–71.

Mason MT (1986b) On the scope of quasi-static pushing. In: *International Symposium on Robotics Research*. Cambridge, MA: MIT Press, pp. 229–233.

Mason MT (2001) *Mechanics of Robotic Manipulation*. Cambridge, MA: MIT Press.

Miller AT, Knoop S, Christensen HI and Allen PK (2003) Automatic grasp planning using shape primitives. In: *Proceedings 2003 IEEE International Conference on Robotics and Automation (ICRA'03)*, vol. 2. IEEE, pp. 1824–1829.

Moreau JJ (1988) Unilateral contact and dry friction in finite freedom dynamics. In: *Nonsmooth Mechanics and Applications*. New York: Springer, pp. 1–82.

Omrcen D, Boge C, Asfour T, Ude A and Dillmann R (2009) Autonomous acquisition of pushing actions to support object grasping with a humanoid robot. In: *Humanoids 2009*.

Parrilo PA (2000) *Structured Semidefinite Programs and Semialgebraic Geometry Methods in Robustness and Optimization*. PhD Thesis, California Institute of Technology.

Peshkin MA and Sanderson AC (1988a) The motion of a pushed, sliding workpiece. *IEEE Journal on Robotics and Automation* 4(6): 569–598.

Peshkin MA and Sanderson AC (1988b) Planning robotic manipulation strategies for workpieces that slide. *IEEE Journal on Robotics and Automation* 4(5): 524–531.

ROBOTIQ (2017) Two-finger adaptive robot gripper. http://robotiq.com/products/adaptive-robot-gripper/ (accessed 15 February 2017).

Stewart DE and Trinkle JC (1996) An implicit time-stepping scheme for rigid body dynamics with inelastic collisions and coulomb friction. *International Journal for Numerical Methods in Engineering* 39(15): 2673–2691.

Whitney DE (1983) Quasi-static assembly of compliantly supported rigid parts. *ASME Journal of Dynamic Systems, Measurement, and Control* 104: 65–77.

Wishart J (1928) The generalised product moment distribution in samples from a normal multivariate population. *Biometrika* 20A (1/2): 32–52.

Yoshikawa T and Kurisu M (1991) Identification of the center of friction from pushing an object by a mobile robot. In: *Proceedings 1991 IEEE/RSJ International Workshop on Intelligence for Mechanical Systems (IROS'91)*. IEEE, pp. 449–454.

Yu KT, Bauza M, Fazeli N and Rodriguez A (2016) More than a million ways to be pushed. a high-fidelity experimental dataset of planar pushing. In: *2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, pp. 30–37.

Zhou J, Bagnell JA and Mason MT (2017) A fast stochastic contact model for planar pushing and grasping: theory and experimental validation. In *Robotics: Science and Systems XIII*.

Zhou J, Paolini R, Bagnell JA and Mason MT (2016) A convex polynomial force-motion model for planar sliding: identification and application. In: *2016 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 372–377.

Zhou J, Paolini R, Johnson AM, Bagnell JA and Mason MT (2017) A probabilistic planning framework for planar grasping under uncertainty. *IEEE Robotics and Automation Letters* 2(4): 2111–2118.

Appendix. Index to multimedia extensions

Archives of IJRR multimedia extensions published prior to 2014 can be found at http://www.ijrr.org, after 2014 all videos are available on the IJRR YouTube channel at http://www.youtube.com/user/ijrrmultimedia

Table of Multimedia Extension

Extension	Media type	Description
1	Video	Model identification and visualization
2	Video	Model applications